

---

# Convex Optimization For Non-Convex Problems via Column Generation

---

**Julian Yarkony**  
 Experian Data Lab  
 San Diego CA  
 julian.e.yarkony@gmail.com

**Kamalika Chaudhuri**  
 University of California San Diego  
 San Diego CA  
 kamalika@cs.ucsd.edu

## Abstract

We apply column generation to approximating complex structured objects via a set of primitive structured objects under either the cross entropy or L2 loss. We use L1 regularization to encourage the use of few structured primitive objects. We attack approximation using convex optimization over an infinite number of variables each corresponding to a primitive structured object that are generated on demand by easy inference in the Lagrangian dual. We apply our approach to producing low rank approximations to large 3-way tensors.

## 1 Introduction

We consider the general problem of approximating a complex structured object using a non-negative weighted combination of primitive structured objects under regularization [10, 11, 4]. Given a class of primitive structured objects (primitives) the set of complex objects is any weighted combination of the primitives or any weighted combination lying in a convex hull or cone [13]. Such problems are common throughout the machine learning literature and include, sparse coding [39], matrix approximation [44, 47], tensor approximation [22], mixture modeling, autoencoder modeling [4], etc. For example in 3-way tensor (three dimensional) approximation [2, 20, 26, 3] the set of rank-1, 3-way tensors are the primitives; the positive cone over rank-1 3-way tensors is the space optimized over.

In this paper we study two families of complex structured object approximation problems. The first (Family One) considers minimizing the L2 loss between a complex structured object versus its approximation with L1 regularization [18] encouraging the use of fewer primitives in the approximation. One example of a problem under Family One is lossy compressing an image so that the colors in the uncompressed image are similar in an L2 sense to the original [47].

The second family (Family Two) considers minimizing the cross entropy loss between a complex structured object corresponding to a probability distribution versus its approximation with L1 regularization encouraging a fewer number of components in the approximation. Sparsity in the L1 sense is achieved by encouraging more of the distribution to be explained by the uniform distribution or other white noise, or high entropy distribution. One example of a problem under Family Two is fitting a mixture of gaussians to describe a probability distribution. Here the high entropy distribution could be the maximum likelihood gaussian fit to all of the data.

The two families differ in the difficulty of approximation. Approximation in Family One is easier than approximation in Family Two, and thus Family One can be applied in larger scale applications. However the cross entropy loss of Family Two is a far more appropriate loss function in domains where probability distributions are being studied, which is common in machine learning.

A key difficulty in complex structured object approximation is that the number of primitive structured objects may be infinite or at least exponential and prevent the construction of a quality approxi-

mation. Much previous work in machine learning relies on non-convex methods such as Expectation Maximization (EM)[17, 43, 5, 48] which make greedy local moves to improve their approximation. Gradient descent also has great practical value especially in the domain of deep neural networks [45, 6, 42, 35]

An alternative approach to solving optimization problems is column generation [21, 7, 51]. Column generation is a powerful method generally used for massive scale integer programming problems originating in operations research such routing flights for large airlines. In this line of work one optimizes over the entire space of exponential number of variables in continuous space. The corresponding LP relaxation is a convex optimization and given a finite number of variables can be solved via interior point methods[13, 31] in polynomial time. However the set of all variables can not be enumerated much less used in optimization. Column generation operates by solving the optimization problem given a small subset of the variables. Next one or more variables are identified that will improve the objective if added to the subset under consideration and these are added to the subset under consideration. The optimization problem is then resolved. This continues until convergence.

In many optimization problems analysis of Lagrange multipliers allows for the exact or near exact computation of the best variable to add to the subset under consideration. Such analysis consists of solving a dramatically simpler version of the original problem which is often polynomial time solvable, or at least can be approximated to high accuracy in polynomial time perhaps with guarantees. In this paper we apply column generation to complex structure approximation in such a way as to circumvent many concerns pertaining to local optima. To our knowledge we are the first to apply column generation to structured object approximation.

## 1.1 Outline

We now provide an outline of this paper. In Section 2 we formally introduce our two families of problems. Next in Section 3 we formulate the problems in terms of convex optimization over enormous or possibly infinite spaces of variables each corresponding to a primitive. Then in Section 4 we formulate optimization in the form of column generation where variables are generated as needed so as to make convex optimization feasible. In Section 5 we show how to identify primitives to add to consideration to improve the objective. We show examples derivations in the domain of tensors (Sections 5.1,5.2) and gaussian mixture models (Section 5.3).

In Section 6 we show experimental results on tensor problems. We show as a function of time and iteration the value of results from optimization for various size synthetic symmetric 3-way tensor problems. We use examples from both Family One and Two. In Section 7 we briefly discuss relevant papers and their relationship to our work. In Section 8 we discuss future work and extensions.

## 2 Formal Model

We now formally discuss Family One and Two. Consider a structured object consisting of a finite number of values denoted  $\bar{T}$ . Here  $\bar{T}$  may correspond to a tensor or a probability distribution for example. We denote the vectorization of  $\bar{T}$  by column vector  $T$ . We denote the set of primitives that may be used to construct  $T$  as  $\mathcal{M}$  and reference its members with  $m$ . Primitives may be rank-1 tensors or gaussian distributions in the previous examples. We define a matrix  $M$  by horizontally stacking the column vectors  $m \in \mathcal{M}$  together. We define a non-negative weighted combination of the columns of  $M$  using a non-negative column vector  $w$  of length equal to  $|\mathcal{M}|$ .

In this document we study the two families in parallel since the approaches for approximation introduced in this paper are highly related. Also much of the notation used to discuss the two families is shared. Similarly many problems that are tackled using Family One have corresponding problems in Family Two.

### 2.1 Family One

We define the optimal model according to L2 loss under L1 regularization as follows. We use  $t$  to denote transpose. We use  $\vec{\ell}$  to define a positive constant column vector of length  $|\mathcal{M}|$  where the constant value is  $\ell \in \mathcal{R}_+$ .

$$\min_{w \geq 0} \frac{1}{2} (T - Mw)^t (T - Mw) + \vec{\ell}^t w \quad (1)$$

## 2.2 Family Two

We now define the optimal model according to cross entropy loss under L1 regularization as follows. We use  $\log$  to denote the element-wise logarithm. We use  $\vec{\ell}_a$  to define a positive constant column vector of length  $|\mathcal{M}|$  where the constant value is  $\ell_a \in \mathcal{R}_+$ . In addition  $\vec{\ell}_a$  that has a single zero valued entry. This entry is associated with a special column  $m_0$  which corresponds to the uniform distribution or other white noise distribution. The entry of  $w$  corresponding to  $m_0$  is denoted  $w_0$ . This is useful in modeling probability distributions where  $T$  is a complex probability distribution and each  $m$  is a primitive probability distribution. The corresponding optimization problem is written below.

$$\min_{\substack{w \geq 0 \\ w^t \mathbf{1} = 1}} -T^t \log(Mw) + \vec{\ell}_a^t w \quad (2)$$

Since the Eq 2 is non-decreasing in  $w_0$  we write Eq 2 as follows.

$$\text{Eq 2} = \min_{\substack{w \geq 0 \\ w^t \mathbf{1} \leq 1}} -T^t \log(Mw) + \vec{\ell}_a^t w \quad (3)$$

## 3 Formulating Optimization

We now consider the treatment of structured object approximation using the tools of quadratic programming (QP) for Family One and linear programming (LP) for Family Two.

### 3.1 Family One

We now introduce a surrogate object  $K$  which corresponds to element-wise distance between  $T$  and  $Mw$ . We now write optimization as a quadratic program.

$$\begin{aligned} \text{Eq 1} &= \min_{\substack{w \geq 0 \\ K \geq 0}} \frac{1}{2} K^t K + \vec{\ell}^t w \\ \text{s.t. } &K \geq T - Mw \\ &K \geq Mw - T \end{aligned} \quad (4)$$

Here Eq 4 corresponds to a convex quadratic program with an intractable  $|\mathcal{M}|$  number of variables (one for each primitive) and a small number of constraints equal to  $2|T|$ . We refer to Eq 4 as the primal problem for Family One and it is associated with a dual problem, which is also a convex quadratic program. This dual problem is written below using dual variables  $\psi, \lambda$  each of which are of length equal to  $|T|$ .

$$\begin{aligned} \text{Eq 1} &= \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} \frac{-\lambda^t \lambda}{2} - \psi^t \lambda - \frac{\psi^t \psi}{2} + T^t \lambda - T^t \psi \\ &\vec{\ell} \geq M^t \lambda - M^t \psi \end{aligned} \quad (5)$$

The dual problem above is derived in Appendix A.

### 3.2 Family Two

We now introduce a surrogate object  $K$  which corresponds to the element-wise logarithm of  $Mw$ .

$$\begin{aligned} \text{Eq 2} &= \min_{\substack{w \geq 0 \\ 1^t w \leq 1 \\ K \geq 0}} T^t K + \vec{\ell}_a^t w \\ \text{s.t. } &-K \leq \log(Mw) \end{aligned} \quad (6)$$

Since the log function is concave we express it as the lower envelope of a set of affine upper bounds each of which is constructed via a first order Taylor expansion. We write this formally below.

$$\log(y) = \min_{\eta \in (0, \infty)} \log(\eta) + (y - \eta) \frac{1}{\eta} \quad \forall y \in (0, \infty) \quad (7)$$

We now apply the lower envelope expression of log in Eq 7 to Eq 2. This produces the linear program below.

$$\begin{aligned} \text{Eq 2} = & \min_{\substack{w \geq 0 \\ 1^t w \leq 1 \\ K \geq 0}} T^t K + \vec{\ell}_a^t w \\ \text{s.t. } & -K \leq \log(1\eta) + \frac{(Mw - 1\eta)}{\eta} \quad \forall \eta \in (0, 1] \end{aligned} \quad (8)$$

We refer to Eq 8 as the primal problem for Family Two. The dual form of Eq 8 which is also a linear program, is written below using dual variables  $\alpha$  and  $\beta$ . Here  $\alpha$  is a scalar and  $\beta$  is associated with a unique vector of cardinality  $|T|$  for every  $\eta$ .

$$\begin{aligned} \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} & -\alpha + \sum_{\eta \in (0, 1]} (1 - 1 \log(\eta))^t \beta^\eta \\ & T \geq \sum_{\eta \in (0, 1]} \beta^\eta \\ & \vec{\ell}_a + 1\alpha \geq M^t \sum_{\eta \in (0, 1]} \frac{1}{\eta} \beta^\eta \end{aligned} \quad (9)$$

The dual problem above is derived in Appendix B.

## 4 Inference in General Terms

Solving a quadratic or linear program is done by various methods such as interior points or simplex. However to employ them there must be a finite number of variables and constraints and our problems do not satisfy this criteria. To circumvent it we solve altered versions of the problems that consider only a subset of the constraints and variables. Variables or constraints are then added when and if they are violated or would improve the objective. We iterate between solving the altered problems and adding variables or constraints. The dual problems facilitate the addition of primal variables via analysis of the dual variables.

### 4.1 Family One

We solve problems in Family One by solving the dual problem. In order to solve the dual problem we must identify a subset of the constraints such that when enforced no other constraints are violated. We build that subset denoted  $\hat{M}$ , which is called the working set, greedily. Here  $\hat{M}$  is initialized to the empty set or with any subset of the columns of  $M$ . We write the corresponding quadratic program below.

$$\begin{aligned} \text{Eq 1} \leq & \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} \frac{-\lambda^t \lambda}{2} - \psi^t \lambda - \frac{\psi^t \psi}{2} + T^t \lambda - T^t \psi \\ & \vec{\ell} \geq \hat{M}^t \lambda - \hat{M}^t \psi \end{aligned} \quad (10)$$

Finding the most violated constraint corresponds to selecting a column of  $M$  to maximize the following for any  $\lambda$  and  $\psi$ . For ease of notation we define a term  $\theta = \lambda - \psi$ .

$$\max_{m \in \mathcal{M}} m^t(\lambda - \psi^t) = \max_{m \in \mathcal{M}} \theta^t m \quad (11)$$

Consider a setting where we are able to identify the maximizing argument  $m$  for Eq 11. The corresponding column generation algorithm is written in Alg 1.

---

**Algorithm 1** Dual Optimization

---

```

 $\hat{M} \leftarrow \{\}$ 

while True do
   $[\lambda, \psi] \leftarrow$  Solve Eq 10 given  $\hat{M}$ 
   $\theta \leftarrow \lambda - \psi$ 
   $[m] \leftarrow \max_{m \in \mathcal{M}} \theta^t m$ 
  if  $\theta^t m > \ell$  then
     $\hat{M} \leftarrow [\hat{M}, m]$ 
  else
    BREAK
  end if
end while

```

---

At the termination of Alg 1 it is the case that Eq 10= Eq 1. Also at any time the primal variables  $w$  which are produced concurrently with  $\lambda, \psi$  describe a valid sub-optimal solution to Eq 1 when using interior points methods. At termination additional sparsity can be created by running least angle regression given the basis  $\hat{M}$  [18].

## 4.2 Family Two

Applying column generation to Family Two is challenging because there are an infinite number of variables in both the primal and dual. To meet this difficulty we store a subset of the constraints/variables to the primal and dual. We then produce a solution using only those constraints/variables. We then identify those violated constraints and add them to the constraint set. Constraints added in the primal correspond to new variables in the dual and similarly constraints added in the dual correspond to new variables in the primal.

To assist our discussion we index  $T$  as follows. We index  $T$  with  $x$  where  $T_x$  is the  $x$ 'th value of vector  $T$ . We index  $K$  and  $Mw$  similarly. We use  $\hat{\mathcal{S}}$  to denote the working subset of  $x, \eta$  pairs and use  $\hat{M}$  to denote the working subset of  $M$ . We initialize  $\hat{M}$  with the uniform distribution/white noise/high entropy distribution  $m_0$  and initialize  $\hat{\mathcal{S}}$  to include  $[x, \frac{1}{|T|}]$  for all  $x$ . The initial setting of  $\hat{\mathcal{S}}$  is an arbitrary choice that worked well in practice. We use  $m$  to index the columns of  $M$ . The corresponding primal and dual pair are below.

$$\begin{aligned} & \min_{\substack{w \geq 0 \\ 1^t w \leq 1 \\ K \geq 0}} T^t K + \vec{\ell}_a^t w & (12) \\ \text{s.t. } & -K_x \leq \log(\eta) + \frac{1}{\eta} (\hat{M}w)_x - 1 \quad \forall [x, \eta] \in \hat{\mathcal{S}} \end{aligned}$$

$$\begin{aligned} \text{Eq 12} = & \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} -\alpha + \sum_{x, \eta \in \hat{\mathcal{S}}} (1 - \log(\eta)) \beta_x^\eta & (13) \\ & T_x \geq \sum_{\substack{[x, \eta] \in \hat{\mathcal{S}} \\ \text{s.t. } \hat{x} = x}} \beta_x^\eta \quad \forall x \\ & \vec{\ell}_a + 1\alpha \geq \sum_{x, \eta \in \hat{\mathcal{S}}} m_x \frac{1}{\eta} \beta_x^\eta \quad \forall m \in \hat{M} \end{aligned}$$

For ease of notation we introduce a vector  $\theta$  defined as follows.

$$\theta_x = \sum_{\substack{[x, \eta] \in \hat{S} \\ s.t. \hat{x} = x}} \frac{1}{\eta} \beta_x^\eta \quad \forall x \quad (14)$$

Consider that we are able to identify violated constraints in the primal and dual. We then apply the following iteration. We solve Eq 13 which provides us with a solution to Eq 12 as well as Eq 13. We then identify violated constraints in the primal and dual. We then add those to the working sets  $\hat{S}$  and  $\hat{M}$ .

Finding constraints that are violated in the primal is trivial. For each  $K_x$  we minimize with respect to  $\eta$  the following  $K_x + \frac{1}{\eta}(\hat{M}w)_x + \log(\eta) - 1$ . We take the derivative with respect to  $\eta$  then set the derivative equal to zero and finally solve for  $\eta$ .

$$\begin{aligned} 0 &= \frac{-1}{\eta^2}(\hat{M}w)_x + \frac{1}{\eta} \\ (\hat{M}w)_x &= \eta \end{aligned} \quad (15)$$

We then add the pair  $x, \eta$  to  $\hat{S}$  if the constraint is violated. We write the corresponding optimization algorithm for Family Two in Alg 2. As in optimization in Family One, additional sparsity can be

---

**Algorithm 2** Dual Optimization

---

```

 $\hat{M} \leftarrow \{m_0\}$ 

 $\hat{S} \leftarrow \{x, \frac{1}{|T|}\} \forall x$ 

while True do
   $[\alpha, \beta, w, K] \leftarrow$  Solve Eq 12/13 given  $\hat{M}$  and  $\hat{S}$ 
  for  $x$  do
     $\hat{S} \leftarrow \hat{S} \cup (x, (\hat{M}w)_x)$ 
  end for
   $[m] \leftarrow \max_{m \in \mathcal{M}} m^t \theta$ 
  if  $\theta^t m > \ell_\alpha + \alpha$  then
     $\hat{M} \leftarrow [\hat{M}, m]$ 
  end if
  if No constraints added this round then
    BREAK
  end if
end while

```

---

created by running least angle regression given  $\hat{M}$  [18].

## 5 Identifying the Most Violating $m$ or Highly Violated $m$

The previous sections reference solving for the optimal  $m$  as the following optimization  $\max_{m \in \mathcal{M}} \theta^t m$ . The difficulty of solving for the optimal  $m$  is problem and problem instance specific. However these problems tend to be relatively easy.

Consider the problem of approximating a high rank 3-way tensor with a low rank tensor under L2 or cross entropy loss and L1 regularization. In this section we demonstrate that solving for  $m$  fits the optimal rank-1 tensor to a tensor described by the dual variables.

In the case of gaussian mixture models we demonstrate that solving for  $\max_{m \in \mathcal{M}} \theta^t m$  corresponds to training the maximum likelihood gaussian under a weighting of the points described by  $\theta$  which can be solved exactly in closed form. In fact  $\theta$  is an unnormalized probability distribution. Here and in all Family Two examples  $\theta$  must be normalized before applying inference and shifting the scalar outside of the max in  $\max_{m \in \mathcal{M}} \theta^t m$ . Normalization is done by dividing each element of  $\theta$  by  $1^t \theta$ .

It should be observed that local optima are not very problematic when identifying  $m$ . One simply needs to find a local optima that has objective value greater than  $\ell$  or  $\ell_a + \alpha$  should one exist. This results in Alg 1/ 2 changing the dual variables and hence  $\theta$  and allow for one to try again to find a good local optimum on a different and perhaps easier problem. Finding a poor local optimum never results in an increase the objective for either Family One or Two. If a poor quality  $m$  is included early in the column generation process it is given zero weight  $w$  by the optimizer. Multiple local optima each corresponding to a violated constraint can be computed here and added to  $\hat{M}$ .

We consider three examples of finding the most violated  $m$  below and with additional examples in Appendix C,D.

### 5.1 Example, Family One: Symmetric 3-way Tensors

We now study  $\max_{m \in M} \theta^t m$  in the domain where  $m$  corresponds to fitting a symmetric 3-way tensor. Consider the case that  $T$  corresponds to the vectorization of a symmetric 3-way tensor. Thus each column of  $M$  also corresponds to a symmetric 3-way tensor. We describe  $m$  using a vector  $v$ . Here  $v \in \mathcal{R}(|T|^{\frac{1}{3}})$  where  $v$  is a unit vector. The non-vectorized form of  $m$  is denoted  $\bar{m}$  and is indexed by  $i, j, k \in \{0, 1, 2 \dots N - 1\}$ . We define  $\bar{m}$  in terms of  $v$  below.

$$\bar{m}_{ijk} = v_i v_j v_k \quad (16)$$

We write optimization below.

$$\max_{m \in M} \theta^t m = \max_{\substack{v \\ v^t v = 1}} \sum_{ijk} \theta_{ijk} v_i v_j v_k \quad (17)$$

The projected gradient update for  $v_i$  for all  $i$  is written below.

$$\begin{aligned} \dot{v}_i &\leftarrow v_i + (\text{stepsize}) \cdot \sum_{jk} \theta_{ijk} v_j v_k \quad \forall i \\ v &\leftarrow \frac{\dot{v}}{\dot{v}^t \dot{v}} \end{aligned} \quad (18)$$

Remarkably the seminal power iteration [37, 40] can be applied in place of projected gradient descent. Furthermore for tensors convergence of the power iteration to a local optima of the objective can be guaranteed [2]. Unlike in the case of 2-way tensor (matrix) global optimality is not guaranteed. The corresponding updates are written below.

$$\begin{aligned} \dot{v}_i &\leftarrow \sum_{jk} \theta_{ijk} v_j v_k \\ v &\leftarrow \frac{\dot{v}}{\dot{v}^t \dot{v}} \end{aligned} \quad (19)$$

We repeatedly update  $v$  until convergence. At termination we add  $m$  and  $-m$  to the working set  $\hat{M}$  when applying the power iteration. This is because the power iteration maximizes the magnitude of  $\theta^t m$  without concern for the sign.

#### 5.1.1 Note on Even way Tensors and Optimization

For even way tensors (2-way,4-way,6-way...) an outer product does not produce all possible rank-1 tensors for  $M$ . It fails to create those constructed by an outer product then multiplied by  $-1$ . This can be seen by observing the following: multiplying the vector  $v$  by  $-1$  does not flip the sign of all elements of  $m$  (it flips none) while this is achieved for odd way tensors (3-way, 5-way,7-way...). Thus at termination of optimization power iteration optimization we add  $m$  and  $-m$  to the working set  $\hat{M}$ .

Projected gradient optimization must be similarly altered in the case of even way tensors. This is done by computing  $\max_{m \in M} \theta^t m$  and  $\max_{m \in M} (-\theta)^t m$  and adding  $m, -m$  to  $\hat{M}$  corresponding to violated constraints.

This is important because Family One is restricted to have non-negative components  $w$  to fit in the standard form for quadratic programming however there need not be a model constraint that the weights are non-negative.

## 5.2 Example Family Two: Fitting a symmetric 3-way tensor defined by probability distribution

We now study  $\max_{m \in M} \theta^t m$  in the domain where  $m$  corresponds to fitting a symmetric 3-way tensor in the setting of Family Two. We use the notation of Section 5.1. In this section we solve optimization relying on Jenson's inequality [27, 28] to force  $m$  to correspond to a probability distribution at all times and avoid gradient descent. We write optimization below.

$$\max_{\substack{v \geq 0 \\ 1^t v = 1}} \sum_{ijk} \theta_{ijk} v_i v_j v_k \quad (20)$$

Recall that  $\theta$  is non-negative and we normalize  $\theta$  to sum to one and shift the normalization constant outside of the max. We now write optimization over the log of  $\max_{m \in M} \theta^t m$ .

$$\max_{\substack{v \geq 0 \\ 1^t v = 1}} \log \left( \sum_{ijk} \theta_{ijk} v_i v_j v_k \right) \quad (21)$$

As in EM methods for inference in probabilistic models we define a proposal probability distribution  $z$  indexed by  $ijk$ . We initialize  $z$  as follows to reflect the probability distribution  $\theta$  though this initialization heuristic and random initialization of  $z$  is also valid.

$$z_{ijk} \leftarrow \frac{\theta_{ijk}}{\sum_{ijk} \theta_{ijk}} \quad (22)$$

We now multiply and divide by  $z$  as is standard in EM methods and apply Jenson's inequality .

$$\begin{aligned} \text{Eq 21} &= \max_{\substack{v \geq 0 \\ 1^t v = 1}} \log \left( \sum_{ijk} \frac{z_{ijk}}{z_{ijk}} \theta_{ijk} v_i v_j v_k \right) \\ &\geq \max_{\substack{v \geq 0 \\ 1^t v = 1}} \sum_{ijk} -z_{ijk} \log z_{ijk} + z_{ijk} \log \theta_{ijk} \\ &\quad + \sum_{ijk} z_{ijk} \log v_i + \sum_{ijk} z_{ijk} \log v_j + \sum_{ijk} z_{ijk} \log v_k \end{aligned} \quad (23)$$

We now add a Lagrange multiplier  $\tau$  to enforce that  $v$  sums to one. There will be no need to enforce non-negativity in optimization.

$$\begin{aligned} \max_{v \geq 0} \min_{\tau \in (-\infty, \infty)} &\tau(1 - 1^t v) + \sum_{ijk} -z_{ijk} \log z_{ijk} \\ &+ \sum_{ijk} z_{ijk} \log \theta_{ijk} + \sum_{ijk} z_{ijk} \log v_i \\ &+ \sum_{ijk} z_{ijk} \log v_j + \sum_{ijk} z_{ijk} \log v_k \end{aligned} \quad (24)$$

We now write the optimization with respect to  $v_i$ . We now take derivative with respect to  $v_i$  and set it equal to 0. Notice that  $v_i$  is present in six terms in the derivative.



$$\begin{aligned}
0 &= -\tau + \sum_{jk} (z_{ijk} + z_{ikj} + z_{jik} + z_{kij} + z_{jki} + z_{kji}) \frac{1}{v_i} \\
\tau v_i &= \sum_{jk} (z_{ijk} + z_{ikj} + z_{jik} + z_{kij} + z_{jki} + z_{kji}) \\
v_i &= \frac{1}{\tau} \sum_{jk} (z_{ijk} + z_{ikj} + z_{jik} + z_{kij} + z_{jki} + z_{kji})
\end{aligned} \tag{25}$$

Observe that  $v_i \propto \sum_{jk} z_{ijk}$ . Since it is the case that  $1^t v = 1$ . Then the following is true.

$$\tau = \sum_{ijk} z_{ijk} \tag{26}$$

The optimizing updates for  $z_{ijk}$  set  $z_{ijk}$  proportional to  $\theta_{ijk} v_i v_j v_k$  based on the standard application of the tightest bound for Jenson's inequality. Therefore the final updates are as follows.

$$\begin{aligned}
z_{ijk} &\propto \theta_{ijk} v_i v_j v_k \\
v_i &\propto \sum_{jk} z_{ijk}
\end{aligned} \tag{27}$$

We repeatedly update  $z$  then  $v$  until convergence.

### 5.3 Example Gaussian Mixture models

Consider that  $T$  represents samples drawn from a continuous probability distribution. Here  $T$  has one index for every one of  $N$  samples. Consider the problem of approximating  $T$  using a set of basis functions  $\mathcal{M}$ . Specifically consider the case of a gaussian basis on one dimension with fixed variance  $\sigma$ . Each  $m$  describes a particular gaussian via displaying the density at each data point  $x$ .

Let us define each  $m$  via a unique mean  $\mu$ . We use  $p_x$  to denote the spatial position of a given point  $x$ . We apply optimization via Jenson's inequality as in Section 5.2 and using the corresponding notation. Consider a probability distribution  $z$  indexed by  $x$ . We now write the objective for optimization and apply Jenson's inequality [28].

Recall that  $\theta$  is non-negative and we normalize  $\theta$  to sum to one and shift the normalization constant outside of the max. We now write optimization over the log of  $\max_{m \in M} \theta^t m$ .

$$\begin{aligned}
\max_{m \in M} \log(\theta^t m) &= \max_{m \in M} \log \sum_x (\theta_x m_x) \\
&= \max_{\mu} \log \sum_x \theta_x \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(p_x - \mu)^2}{2\sigma^2}} \\
&= \max_{\mu} \log \left( \sum_x \frac{z_x}{z_x} \theta_x \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(p_x - \mu)^2}{2\sigma^2}} \right) \\
&\geq \max_{\mu} \sum_x -z_x \log z_x \\
&\quad + z_x \log \left( \frac{1}{\sigma \sqrt{2\pi}} \theta_x e^{-\frac{(p_x - \mu)^2}{2\sigma^2}} \right) \quad \forall [z \geq 0, 1^t z = 1]
\end{aligned} \tag{28}$$

Given  $z$  we can optimize with respect to  $\mu$ . Optimizing with respect to  $z$  is done as is standard in Jenson's inequality and is written below.

$$z_x \propto \theta_x \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(p_x - \mu)^2}{2\sigma^2}} \tag{30}$$

Similarly given  $\mu$  we can optimize  $z$ . We now write optimization over  $\mu$  given  $z$ .

$$\begin{aligned}
& \max_{\mu} \sum_x -z_x \log z_x \\
& + z_x \log \left( \frac{1}{\sigma \sqrt{2\pi}} \theta_x e^{-\frac{(p_x - \mu)^2}{2\sigma^2}} \right) \\
= & \max_{\mu} \sum_x -z_x \log z_x - z_x \log \sigma - z_x \frac{1}{2} \log(2\pi) \\
& + z_x \log \theta_x + z_x \frac{-(p_x - \mu)^2}{2\sigma^2}
\end{aligned} \tag{31}$$

We now take the derivative with respect to  $\mu$  and set it equal to zero and thus we obtain a closed form expression of  $\mu$ .

$$\begin{aligned}
0 &= \frac{2}{2\sigma^2} \sum_x z_x (\mu - p_x) \\
\sum_x z_x p_x &= \mu \sum_x z_x \\
\sum_x z_x p_x &= \mu
\end{aligned} \tag{32}$$

We repeatedly update  $z$  then  $\mu$  until convergence. We begin with a value for  $\mu$  then solve for  $z$  where we initialize  $\mu$  to correspond to a point  $p_x$  where  $x$  is selected with probability proportionate to  $\theta_x$ . This is an initialization heuristic but attempts to place the gaussian in a region of high density of probability mass  $\theta$ .

## 6 Experiments

In this section we show the effectiveness of our approach for approximating large symmetric 3-way tensors in Family One and Two.

### 6.1 Family One

We now study approximation to symmetric 3-way tensors for Family One to test the effectiveness of our approach. We use various sizes (30,35,40,45), constructed from a convex combination of three to ten unique rank-1 tensors with weights summing to one. We construct each unique rank-1 tensor by the triple outer product of a unique random unit vector. Each such vector has between six and fifteen non-zero elements. We inject noise describing between one and twenty percent of the tensor. For each problem instance we use four different L1 regularizers [0.1, 0.4, 0.7, 1.0]. We consider 9000 problems instances and we continue optimization for up to five minutes per instance after which termination is done after solving the current QP.

In Fig 1(a) we show the loss with respect to time averaged over non-terminated problem instances. This plot demonstrates that we rapidly produce low cost solutions.

In Fig 1(b) we show a scatter plot of total optimization time vs the loss at termination where each instance is a single data point. We show normalized and un-normalized values where normalization corresponds to subtracting the loss on the ground truth model (note that the ground truth model does not model the noise). This plot demonstrates that we are able to fit the tensors rapidly and are able to overfit which is important for an optimization approach.

In Fig 1(c) we show the derivative of the objective with respect to the columns in ground truth model basis. This plot demonstrates that little is gained by adding the ground truth basis tensors to  $\hat{M}$  if not present at convergence. From this we conclude that the power iteration is able to find good local optima of  $\max_{m \in \mathcal{M}} \theta^t m$ .

## 6.2 Family Two

We test the effectiveness of our approach on symmetric 3-way tensors in Family Two exactly as for Family One though on a smaller scale of problems. We use tensors of size twenty where the basis vectors used to construct the tensor have between four and six non-zero values. Each basis vector was non-negative and its elements summed to one instead of having unit norm as for Family One. We consider 750 problem instances and we continue optimization for up to five minutes per instance after which termination is done after solving the current LP.

In Fig 1(d) we show the loss with respect to time averaged over non-terminated problem instances. This plot demonstrates that we rapidly produce low cost solutions though not as quickly as in Family One.

In Fig 1(e) we show the corresponding plot for Family Two Tensors for 1(b). When normalizing we require the objective of the ground truth model. We add a small amount of probability mass (0.00001) to each entry of the ground truth model so that it does not have any zero values during cross entropy computation. Thus the sum of the values elements in the ground truth model is greater than one. This plot demonstrates that we are able to fit the tensors rapidly and are able to overfit which is important for an optimization approach.

In Fig 1(f) we show the corresponding plot for Family Two Tensors for 1(c). This plot demonstrates that little is gained by adding the ground truth basis tensors to  $\hat{M}$  if not present at convergence. From this we conclude that our Jensen's inequality based optimization is able to find good local optima of  $\max_{m \in \mathcal{M}} \theta^t m$ .

## 7 Literature Review

Our work can be positioned at the intersection of two well studied areas: 1) column generation for combinatorial optimization and 2) efficient representation of data notably with regards to low rank tensor decompositions. In this section we discuss some of the related work in these two vast areas in the context of our work.

### 7.1 Column Generation

#### 7.1.1 The Cutting Stock Problem

Our work is intimately related to the classical work of [21] on the cutting stock problem which originated in the paper industry. In [21] the authors determine how to satisfy a set of demands for rolls of paper of various widths given rolls of paper of longer widths while minimizing the scrap. In this line of work there are a massive number of variables where each variable corresponds to a different way to cut a roll. Each way of cutting a roll is called a pattern. The number of possible patterns can grow exponentially in the number of unique widths demanded.

The following very small scale example illustrates the concept of a pattern. Consider cutting a 100 meter wide roll of paper in the context of demands for 30 and 45 meter long rolls. Two example patterns would be (1) cut the roll into three 30 meter long rolls with 10 meters of scrap (2) cut the roll into one 45 meter long roll and one 30 meter long roll with 25 meters of scrap. Optimization is formulated as an integer linear program (ILP) that is NP-Hard. The value of a variable in a solution to the ILP denotes the number of rolls that are cut using the pattern associated with that variable. The integer linear program is relaxed to a linear program allowing for a fractional number of each pattern to be used. Optimization is initialized with a fixed number of patterns sufficient to satisfy all orders but likely with excess waste. Optimization proceeds by solving the linear program given the current set of patterns, followed by generating new patterns. This is repeated until convergence.

Finding the optimal pattern to add corresponds to solving a knapsack problem[30, 29]. The knapsack problem is NP-Hard however it can be approximated very well using a dynamic programming based polynomial time approximation scheme. This can be understood as maximizing the reward obtained by cutting a single roll by choosing a single pattern given reward values associated with each width of paper. Here the rewards correspond to the value of the dual variables in the dual linear program. In this manner the problem of optimizing over a massive number of patterns is reduced to finding a single pattern.

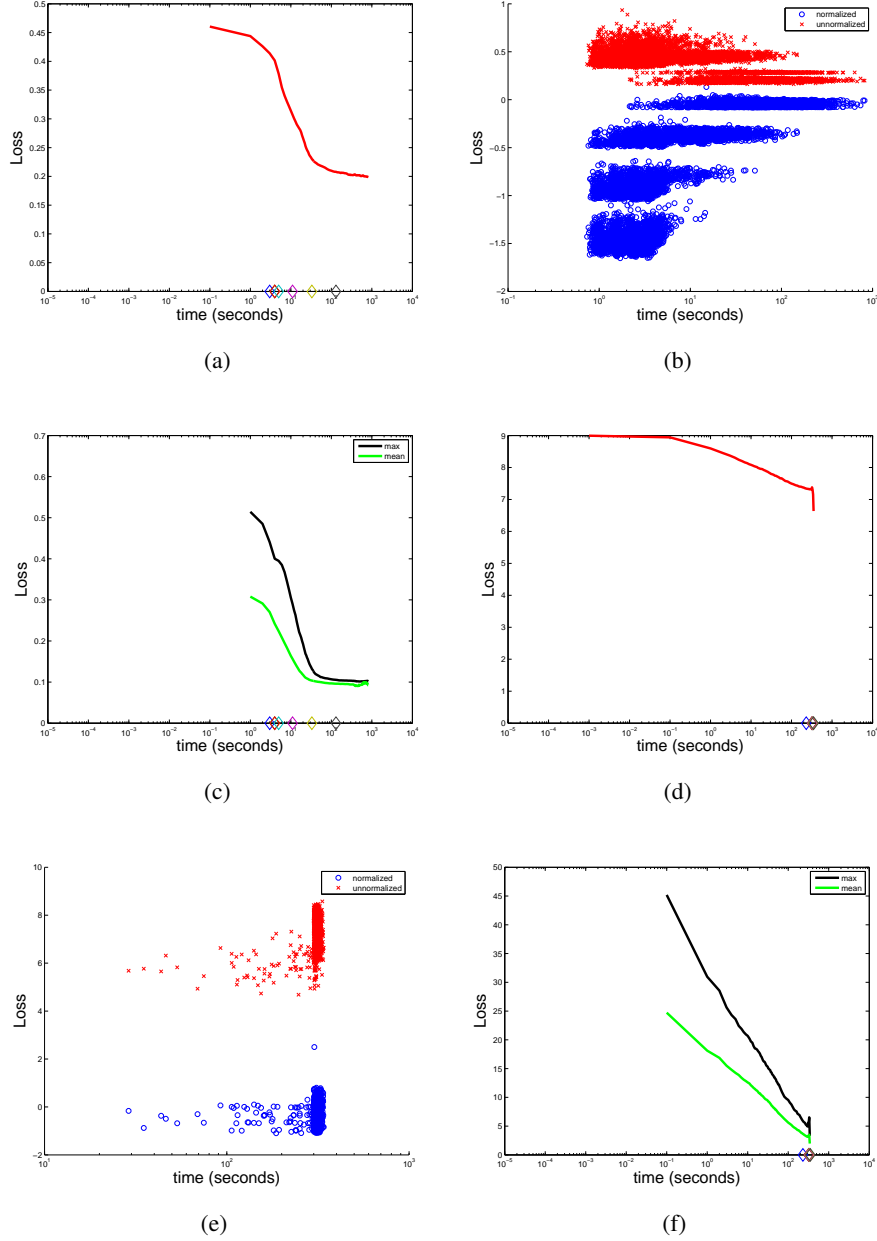


Figure 1: We show quantities averaged over problem instances. 1(a) (Family One) We show loss with respect to time averaged over problem instances that have not terminated. We use diamonds to denote [95,80,65,50,35,20,5] percent not terminated. 1(b) (Family One) We show a scatter plot of total optimization time vs loss at termination of optimization. Each instance is associated with a pair of points, one associated with normalized loss and other other associated with unnormalized loss. Normalization corresponds to subtracting the loss corresponding to the model. 1(c) (Family One) We show the derivative of the objective with respect to the columns in ground truth model basis. We show the maximum and mean value. We average over non-terminated instances. We use diamonds to denote [95,80,65,50,35,20,5] percent non-terminated. Fig1(d),1(e),1(f) are the corresponding plots for Family Two for Fig 1(a),1(b),1(c) respectively.

In our applications we approximate high rank tensors with low rank tensors. When generating new rank-1 tensors to add to our set of primitives (working set) we construct a tensor to maximize sum of the element-wise products between it and another tensor corresponding to the dual variables. In this manner the problem of optimizing over an infinite number of rank-1 tensors is reduced to iteratively fitting a rank-1 tensor. Like the knapsack problem, fitting the optimal rank-1 tensor is known to be NP-Hard [24] however like the knapsack problem well studied optimization schemes exist [32].

Producing integer solutions in the cutting stock may be done heuristically by greedily rounding up fractional values to integer values or by a principled but expensive branch and price algorithm[7]. Fortunately our use of L1 regularization does not require integer solutions so branch and price is not needed. Variants of the cutting stock problem include enforcing a hard constraint on the maximum number of patterns used and this can be interpreted as L0 regularization and is akin to our L1 regularization.

### 7.1.2 Marginals in Graphical Models

Our work is related to the work of [9] (which is extended in [34]). While column generation is not explicitly used or mentioned, their approach is very similar to column generation and their problem has a cross entropy loss like in Family Two. In this paper the authors attack the problem of computing the marginal distributions of variables in a Markov random field (MRF). This is done by minimizing standard Bethe-style convex variational objective. The authors introduce an algorithm to construct a probability distribution over solutions and adds solutions to the probability distribution greedily. At each step of the algorithm a new solution is generated and added to the probability distribution.

Choosing the optimal solution to add corresponds to MAP inference which is an integer program and is often NP-Hard. The potentials for MAP inference are a function of the current probability distribution and the potentials that define the original MRF, unlike the dual variables in our methods. MAP Inference can be challenging because problem structure such as sub-modularity [14] which allows for efficient MAP inference is lost when present in the original MRF however LP relaxation techniques[33, 46] can be used to produce efficient approximate solutions. Once computed a line search is used to determine how much to weight this solution versus the other solutions. The determination of the optimal solution via MAP inference can be interpreted as adding a primal variable (primitive) to the working set while the line search is a greedy move towards optimizing the cross entropy loss.

### 7.1.3 Generating Multiple Primal Variables

Finding the optimal primitive variable to add in our applications is NP-Hard but multiple local optima can be generated concurrently given the same set of dual variables with each optima being computed on a separate CPU. This can be contrasted to the work of [50] where large numbers of variables each of which improves the objective are generated in polynomial time concurrently as part of a common operation. In [50], tracking thousands of objects in video is formulated as an integer program which is relaxed to a linear program where each variable corresponds to an entire track of an object. Finding the optimal variable to add to the working set corresponds to solving a dynamic program which is polynomial time solvable to global optimality. However dynamic programming provides not only the optimal solution but the optimal solution passing through each position in space-time for free. Each of these or a subset of the most violated can be added to the working set. [50] adds one thousand such tracks per iteration allowing for very fast inference. Surprisingly solving the LP is not hindered by the addition of large numbers of tracks by huge working sets. In our approach computing multiple local optima allows for some of the benefits of multiple solutions per dual solution it does not take advantage of the scale for easy generation of variables of [50].

## 7.2 Low Rank Tensor Decompositions

### 7.2.1 Orthogonal Tensor Decompositions

In [2], the authors attack the problem of latent variable modeling leveraging modeling assumptions about how the data was generated. If these assumptions are obeyed then they can guarantee globally optimal recovery of the data generating model. Inference consists of computing a low rank approx-

imation of an orthogonal tensor generated using the triple co-occurrence of moments in the data set (plus whitening). The work of [2] is associated with powerful statistical guarantees. However if the modeling assumptions are not obeyed by the data generation process the recovered solution which is intended to describe a mixture probability of distributions need not produce probability distributions though these can be rounded to probability distributions to produce approximate solutions.

In contrast to the work of [2] we focus on optimization, and we ensure that the outputted representation follows the rules that define the primitives. Like [2] we use the power iteration to generate primitives though only in the context of creating a rank-1 representations of the dual variables.

### 7.3 Gradient Descent and Alternating Least Squares

Two common and powerful approaches for tensor decomposition are gradient descent [1] and alternating least squares (ALS)[16, 32, 23]. Gradient descent optimizes all modes of the tensor concurrently and converges to a local optima of the objective. Alternatively ALS solves for one of the modes of the tensor at a time keeping the other modes fixed. This is achieved via least squares and is coordinate-wise optimal. ALS is not guaranteed to achieve global optimality but in practice is highly successful and benefits from the lack of a need for a step size as in gradient methods[32].

In Appendix C, D we show coordinate-wise updates for Family One and Family Two respectively while gradient descent updates are present in Section 5.1, and in Appendix C. Since we compute only a single, rank-1 tensor at a time our method does not solve least squares problems during coordinate-wise updates. Instead simpler closed form updates exist; for Family One this involve computing a gradient and obtaining an exact solution for the mode.

A major difference between the gradient descent approaches, and ALS in contrast to our work is the form of regularization. For gradient descent and ALS approaches the number of rank-1 terms is fixed in the beginning of optimization while in our approach it is not. Thus gradient descent and ALS have implicit L0 regularization. We have no ability to enforce an L0 norm in our method and must be content to approximate it with an L1 norm.

Another point of difference is that our approach generates new rank-1 tensors but does not adapt those that have already been produced. Gradient descent and ALS continuously alter the tensors they are operating on while ours adds tensors to the working set during optimization leaving the previous ones fixed. Our approach can be used in complementary way to ALS and gradient descent so as to update the rank-1 tensors in the working set perhaps speeding up inference.

#### 7.3.1 Constrained Tensor Decomposition via the Alternating Direction of Multipliers Method

In [36] the authors employ the alternating direction of multipliers method [12] (ADMM) to break down constrained tensor factorization problems into separate problems (sub-problems) that are enforced to have a common solution using Lagrange multipliers. Solutions to the sub-problems are much easier to compute than solving the original constrained problem though they must be re-solved many times. This difficulty can be reduced by leveraging parallel computation.

Lagrange multipliers are used to couple the problems together and operate in such a way as to not make the unconstrained problems more difficult to solve. The work of [36] is extended in [25] where it is demonstrated to produce state of the art results for non-negative tensor factorization, and other problems. Optimization is guaranteed to converge to a stationary point of the constrained objective. Our approach can be used in a complementary way to [36] by using ADMM during the stage of generating new tensors. However unlike in [36, 25] only a single rank-1 tensor would be constructed at a time.

## 8 Conclusions and Future Work

In this document we apply column generation to approximating complex structured objects via a set of primitive structured objects under two families of loss functions with L1 regularization encouraging the use of few structured primitive objects. We attack approximation using convex optimization over an infinite number of variables each of which are generated on demand using the corresponding

dual problem. We apply our approach to producing low rank approximations to large 3-way tensors. Our work provides a broad domain for extensions and we note a few below.

**Increasing Scale:** Solving much larger problems will require clever use of LP and QP solvers. One can attack this by using sub-optimal dual solutions to compute violated constraints. Similarly one can use the LP and QP solvers in ways that do not restart optimization from scratch each time called or which forget constraints that are not active yet slow optimization.

**Diverse Solutions:** Additional speed may be obtained by applying sampling and deterministic sampling approaches such as [41, 49, 8, 34] to add diverse columns to  $\hat{M}$  after each solution to the LP or QP. This would be useful in a domain where solving the QP/LP is significantly more time intensive than maximizing  $\theta^t m$ .

**Size Reduction:** Applications of tensor size reduction methods as used in [2] may prove to be invaluable to extending our work to much larger problems especially if the corresponding statistical guarantees can be preserved. This is challenging as the constraints on elements  $m$  change in the new reduced size space.

**L0 regularization:** The application of branch and bound/price techniques [7] after the construction of  $\hat{M}$  using our approach with L1 regularization, may prove useful for optimizing under L0 regularization.

**Priors and Mixture Modeling:** We did not explore the use of priors over the variance of gaussians for learning mixture of gaussians and optimization over the variance. We suspect this can be accomplished by altering the L1 regularizer to penalize lower variance gaussians and enforcing this via the conjugate prior for the gaussian, which is the normal Wishart prior [38]. This may be useful in the domain of quantum chemistry.

**Sparse Coding/Topic Models** Applying the column generation approach to the Family One problem of sparse coding requires introducing a term to limit the total number of primitives used across documents (data samples). This can be done in an L1 sense by associating each primitive with variable in the QP that is associated with the maximum amount this primitive is used by any document. L1 regularization is then applied to these variables in the QP with the aim of limiting the total number of primitives used across documents. A separate L1 regularization is used to encourage each document to have a sparse representation. The Family Two version of this would be a topic model.

**Deep Relationships:** We suspect that there is a relationship explaining the effectiveness of large single hidden layer neural networks studied in [15], using our work and boosting [19].

## References

- [1] E. Acar, T. G. Kolda, and D. M. Dunlaval. Cppt: Optimization for fitting candecomp/parafac models. *CASTA2008*, page 50, 2008.
- [2] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *J. Mach. Learn. Res.*, 15(1):2773–2832, Jan. 2014.
- [3] A. Anandkumar, D. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden markov models. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, pages 33.1–33.34, 2012.
- [4] P. Baldi. Autoencoders, unsupervised learning, and deep architectures. *Unsupervised and Transfer Learning Challenges in Machine Learning, Volume 7*, page 43, 2012.
- [5] P. Baldi and S. Brunak. *Bioinformatics: The Machine Learning Approach*. MIT Press, Cambridge, MA, USA, 2nd edition, 2001.
- [6] P. Baldi, P. Sadowski, and D. Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5, 2014.
- [7] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46:316–329, 1996.
- [8] D. Batra, P. Yadollahpour, A. Guzman-Rivera, and G. Shakhnarovich. Diverse m-best solutions in markov random fields. In *Computer Vision—ECCV 2012*, pages 1–16. Springer, 2012.
- [9] D. Belanger, D. Sheldon, and A. McCallum. Marginal inference in mrfs using frank-wolfe. In *NIPS Workshop on Greedy Optimization, Frank-Wolfe and Friends*, 2013.
- [10] A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.

- [11] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, Mar. 2003.
- [12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [13] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [14] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:2001, 2001.
- [15] C. Bucilu, R. Caruana, and A. Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541. ACM, 2006.
- [16] J. D. Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of “eckart-young” decomposition. *Psychometrika*, 35(3):283–319, 1970.
- [17] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal Of The Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [18] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.
- [19] Y. Freund and R. Schapire. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- [20] L. Gauvin, A. Panisson, and C. Cattuto. Detecting the community structure and activity patterns of temporal networks: a non-negative tensor factorization approach. *PloS one*, 9(1):e86028, 2014.
- [21] P. Gilmore and R. Gomory. A linear programming approach to the cutting-stock problem. *Operations Research (volume 9)*, 1961.
- [22] L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.
- [23] R. A. Harshman. Foundations of the parafac procedure: Models and conditions for an “explanatory” multi-modal factor analysis. 1970.
- [24] C. J. Hillar and L.-H. Lim. Most tensor problems are np-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.
- [25] K. Huang, N. D. Sidiropoulos, and A. P. Liavas. A flexible and efficient algorithmic framework for constrained matrix and tensor factorization. *arXiv preprint arXiv:1506.04209*, 2015.
- [26] M. Itskov. *Tensor Algebra and Tensor Analysis for Engineers: With Applications to Continuum Mechanics*. Springer Publishing Company, Incorporated, 2nd edition, 2009.
- [27] J. L. W. V. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, 1906.
- [28] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- [29] R. M. Karp. *Reducibility among combinatorial problems*. Springer, 1972.
- [30] H. Kellerer, U. Pferschy, and D. Pisinger. *Introduction to NP-Completeness of knapsack problems*. Springer, 2004.
- [31] K. Koh, S.-J. Kim, and S. P. Boyd. An interior-point method for large-scale l1-regularized logistic regression. *Journal of Machine learning research*, 8(8):1519–1555, 2007.
- [32] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [33] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 28(10):1568–1583, 2006.
- [34] R. G. Krishnan, S. Lacoste-Julien, and D. Sontag. Barrier frank-wolfe for marginal inference. In *Advances in Neural Information Processing Systems*, pages 532–540, 2015.
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [36] A. P. Liavas and N. D. Sidiropoulos. Parallel algorithms for constrained tensor factorization via alternating direction method of multipliers. *Signal Processing, IEEE Transactions on*, 63(20):5450–5463, 2015.
- [37] R. Mises and H. Pollaczek-Geiringer. Praktische verfahren der gleichungsauflösung. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 9(2):152–164, 1929.
- [38] K. P. Murphy. Conjugate bayesian analysis of the gaussian distribution. *def*, 1(2σ2):16, 2007.
- [39] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.
- [40] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999.
- [41] G. Papandreou and A. L. Yuille. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 193–200. IEEE, 2011.



- [42] G. Pollastri, D. Przybylski, B. Rost, and P. Baldi. Improving the prediction of protein secondary structure in three and eight classes using recurrent neural networks and profiles. *Proteins: Structure, Function, and Bioinformatics*, 47(2):228–235, 2002.
- [43] D. Reynolds. Gaussian mixture models. *Encyclopedia of Biometrics*, pages 827–832, 2015.
- [44] S. Roweis. Em algorithms for pca and spca. In *Advances in Neural Information Processing Systems*, pages 626–632. MIT Press, 1998.
- [45] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Technical report, DTIC Document, 1985.
- [46] D. Sontag, T. Meltzer, A. Globerson, T. Jaakkola, and Y. Weiss. Tightening lp relaxations for map using message passing. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 503–510, july 2008.
- [47] M. Turk and A. Pentland. Face recognition using eigenfaces. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR '91., IEEE Computer Society Conference on*, pages 586–591, Jun 1991.
- [48] M. Weber, M. Welling, and P. Perona. Unsupervised learning of models for visual object class recognition, 1999.
- [49] M. Welling. Herding dynamical weights to learn. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1121–1128. ACM, 2009.
- [50] S. Wolf, F. A. Hamprecht, and J. Yarkony. Tracking objects with higher order interactions using delayed column generation. *arXiv preprint arXiv:1512.02413*, 2015.
- [51] J. Yarkony and C. Fowlkes. Planar ultrametrics for image segmentation. In *Neural Information Processing Systems*, 2015.

## A Family One: Dual Derivation

We now derive the dual problem of Family One given the primal problem of Family One. We begin with the Family One objective function after the introduction of the surrogate complex structured object  $K$ .

$$\begin{aligned} \min_{\substack{w \geq 0 \\ K \geq 0}} & \frac{1}{2} K^t K + \ell^t w \\ \text{s.t. } & K \geq T - Mw \\ & K \geq Mw - T \end{aligned} \quad (33)$$

Next we replace the constraints with Lagrange Multipliers.

$$\begin{aligned} \text{Eq 4} = \min_{\substack{w \geq 0 \\ K \geq 0}} \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} & \frac{1}{2} K^t K + \ell^t w \\ & + \psi^t (-T + Mw - K) \\ & + \lambda^t (T - Mw - K) \end{aligned} \quad (34)$$

Since Eq 4 is a convex quadratic program we are able to flip the order of the min and the max without altering the value of the objective.

$$\begin{aligned} \text{Eq 4} = \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} \min_{\substack{w \geq 0 \\ K \geq 0}} & \frac{1}{2} K^t K + \ell^t w \\ & + \psi^t (-T + Mw - K) \\ & + \lambda^t (T - Mw - K) \end{aligned} \quad (35)$$

We now take the first derivative of Eq 35 with respect to  $K$  and then solve for  $K$  in terms of the dual variables  $\lambda$  and  $\psi$ .

$$\begin{aligned} 0 &= K - \psi - \lambda \\ K &= \psi + \lambda \end{aligned} \quad (36)$$

Given the closed form solution for  $K$  in Eq 36 we write Eq 4 with  $K$  replaced by  $\psi + \lambda$ .

$$\begin{aligned} \text{Eq 4} = \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} \min_{w \geq 0} & \frac{1}{2}(\lambda + \psi)^t(\lambda + \psi) + \vec{\ell}^t w \\ & + \psi^t(-T + Mw - \lambda - \psi) \\ & + \lambda^t(T - Mw - \lambda - \psi) \end{aligned} \quad (37)$$

We now group the terms by primal variable.

$$\begin{aligned} \text{Eq 4} = \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} \min_{w \geq 0} & \frac{-1}{2}(\lambda + \psi)^t(\lambda + \psi) + \lambda^t T - \psi^t T \\ & + (\vec{\ell} - \lambda^t M + \psi^t M)w \end{aligned} \quad (38)$$

We now convert the Lagrangian to a convex quadratic program by converting the primal variables  $w$  into constraints.

$$\begin{aligned} \text{Eq 4} = \max_{\substack{\lambda \geq 0 \\ \psi \geq 0}} & \frac{-\lambda^t \lambda}{2} - \psi^t \lambda - \frac{\psi^t \psi}{2} + T^t \lambda - T^t \psi \\ & \vec{\ell} \geq M^t \lambda - M^t \psi \end{aligned} \quad (39)$$

## B Family Two: Dual Derivation

We now derive the dual problem of Family Two given the primal problem of Family Two. We begin with the Family Two objective function after the introduction of the surrogate complex structured object  $K$  and the replacement of log by a concave envelope of affine functions.

$$\begin{aligned} & \min_{\substack{w \geq 0 \\ 1^t w \leq 1 \\ K \geq 0}} T^t K + \vec{\ell}_a^t w \\ \text{s.t. } & -K \leq \log(1\eta) + \frac{(Mw - 1\eta)}{\eta} \quad \forall \eta \in (0, 1] \end{aligned} \quad (40)$$

We now insert Lagrange multipliers  $\alpha$  and  $\beta$ .

$$\begin{aligned} \text{Eq 40} = \min_{\substack{w \geq 0 \\ K \geq 0}} \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} & T^t K + \vec{\ell}_a^t w \\ & \alpha(1^t w - 1) \\ & \sum_{\eta \in (0, 1]} \beta^{t\eta} \left( -K - \frac{1}{\eta} Mw + 1 - 1 \log(\eta) \right) \end{aligned} \quad (41)$$

We now group by primal variable.

$$\begin{aligned} \text{Eq 40} = \min_{\substack{w \geq 0 \\ K \geq 0}} \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} & \sum_{\eta \in (0, 1]} \beta^{t\eta} (1 - 1 \log(\eta)) - \alpha \\ & (T^t - \sum_{\eta \in (0, 1]} \beta^{t\eta}) K \\ & (\vec{\ell}_a^t + \alpha 1^t - \sum_{\eta \in (0, 1]} \frac{1}{\eta} \beta^{t\eta} M) w \end{aligned} \quad (42)$$

We now convert primal variables to hard constraints.

$$\begin{aligned} \text{Eq 40} &= \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} \sum_{\eta \in (0,1]} \beta^{t\eta} (1 - 1 \log(\eta)) - \alpha \\ &\quad T^t - \sum_{\eta \in (0,1]} \beta^{t\eta} \geq 0 \\ &\quad \vec{\ell}_a^t + \alpha 1^t - \sum_{\eta \in (0,1]} \frac{1}{\eta} \beta^{t\eta} M \geq 0 \end{aligned} \quad (43)$$

We now take the transpose of our objective and move terms.

$$\begin{aligned} \text{Eq 40} &= \max_{\substack{\alpha \geq 0 \\ \beta \geq 0}} -\alpha + \sum_{\eta \in (0,1]} (1 - 1 \log(\eta))^t \beta^\eta \\ &\quad T \geq \sum_{\eta \in (0,1]} \beta^\eta \\ &\quad \vec{\ell}_a + 1\alpha \geq M^t \sum_{\eta \in (0,1]} \frac{1}{\eta} \beta^\eta \end{aligned} \quad (44)$$

## C Optimizing $\theta^t m$ , Family One: Non-Symmetric Tensors

We now study  $\max_{m \in M} \theta^t m$  in the domain where  $m$  corresponds to fitting a non-symmetric 3-way tensor. Thus each column of  $M$  also corresponds to a non-symmetric 3-way tensor. We describe  $m$  using vectors unit vectors  $v^a, v^b, v^c$  which are indexed by  $i, j, k$  respectively. The non-vectorized form of  $m$  is denoted  $\bar{m}$  and defined below.

$$\bar{m}_{ijk} = v_i^a v_j^b v_k^c \quad (45)$$

We write optimization below.

$$\max_{m \in M} \theta^t m = \max_{\substack{v^a, v^b, v^c \\ v^{at} v^a = 1 \quad v^{bt} v^b = 1 \\ v^{ct} v^c = 1}} \sum_{ijk} \theta_{ijk} v_i^a v_j^b v_k^c \quad (46)$$

The projected gradient update for  $v_i^a$  for all  $i$  is written below.

$$\begin{aligned} \dot{v}_i^a &\leftarrow v_i^a + \text{stepsize} \sum_{jk} \theta_{ijk} v_j^b v_k^c \quad \forall i \\ v^a &\leftarrow \frac{\dot{v}^a}{\dot{v}^{at} \dot{v}^a} \end{aligned} \quad (47)$$

As an alternative or supplement to gradient descent one can perform coordinate-wise updates which update one of the vectors given the other vectors. Consider optimizing  $v^a$  given  $v^b$  and  $v^c$ . Observe that gradient of  $\sum_{ijk} \theta_{ijk} v_i^a v_j^b v_k^c$  with respect to  $v^a$  has no dependency on  $v^a$  and  $v^a$  is restricted to be unit norm. Thus the optimal  $v^a$  is thus simply the gradient of  $\sum_{ijk} \theta_{ijk} v_i^a v_j^b v_k^c$  with respect to  $v^a$  properly scaled to be a unit vector. Thus we can write the update for  $v^a$  as the following expression.

$$\begin{aligned} \dot{v}_i^a &\leftarrow \sum_{jk} \theta_{ijk} v_j^b v_k^c \quad \forall i \\ v^a &\leftarrow \frac{\dot{v}^a}{\dot{v}^{at} \dot{v}^a} \end{aligned} \quad (48)$$

Optimization over  $v^a, v^b, v^c$  proceeds by cycling through  $a, b, c$  optimizing  $v^a$ , then  $v^b$ , then  $v^c$ . Higher order generalizations of the singular value decomposition may find use here allowing for example  $v^a, v^b$  to be updated given  $v^c$ .

## D Optimizing $\theta^t m$ , Family Two: Non-Symmetric Tensors

We now consider updates for 3-way non-symmetric tensors in Family Two using the notation of Section C. Recall that  $\theta$  is non-negative and that we normalize  $\theta$  to sum to one and shift the normalization constant outside of the max. We now write optimization over the log of  $\max_{m \in M} \theta^t m$ .

$$\max_{\substack{v^a \geq 0, v^b \geq 0, v^c \geq 0 \\ 1^t v^a = 1, 1^t v^b = 1, 1^t v^c = 1}} \log \left( \sum_{ijk} \theta_{ijk} v_i^a v_j^b v_k^c \right) \quad (49)$$

As in EM methods for inference in probabilistic models we define a proposal probability distribution  $z$  indexed by  $ijk$ . We initialize  $z$  as follows to reflect the probability distribution  $\theta$  though this initialization is heuristic and random initialization of  $z$  is also valid.

$$z_{ijk} \leftarrow \frac{\theta_{ijk}}{\sum_{ijk} \theta_{ijk}} \quad (50)$$

We now multiply and divide by  $z$  as is standard in EM methods and apply Jensen's inequality .

$$\begin{aligned} \text{Eq 49} &= \max_{\substack{v^a \geq 0, v^b \geq 0, v^c \geq 0 \\ 1^t v^a = 1, 1^t v^b = 1, 1^t v^c = 1}} \log \left( \sum_{ijk} \frac{z_{ijk}}{z_{ijk}} \theta_{ijk} v_i^a v_j^b v_k^c \right) \quad (51) \\ &\geq \max_{\substack{v^a \geq 0, v^b \geq 0, v^c \geq 0 \\ 1^t v^a = 1, 1^t v^b = 1, 1^t v^c = 1}} \sum_{ijk} -z_{ijk} \log z_{ijk} + z_{ijk} \log \theta_{ijk} \\ &\quad + \sum_{ijk} z_{ijk} \log v_i^a + \sum_{ijk} z_{ijk} \log v_j^b + \sum_{ijk} z_{ijk} \log v_k^c \end{aligned}$$

We now add a Lagrange multiplier  $\tau^a, \tau^b, \tau^c$ , to enforce that  $v^a, v^b, v^c$  sum to one respectively. There will be no need to enforce non-negativity in optimization.

$$\begin{aligned} \max_{\substack{v^a \geq 0 \\ v^b \geq 0 \\ v^c \geq 0}} \min_{\substack{\tau^a \in (-\infty, \infty) \\ \tau^b \in (-\infty, \infty) \\ \tau^c \in (-\infty, \infty)}} &\tau^a (1 - 1^t v^a) + \tau^b (1 - 1^t v^b) + \tau^c (1 - 1^t v^c) + \sum_{ijk} -z_{ijk} \log z_{ijk} \quad (52) \\ &+ z_{ijk} \log \theta_{ijk} + \sum_{ijk} z_{ijk} \log v_i^a + \sum_{ijk} z_{ijk} \log v_j^b + \sum_{ijk} z_{ijk} \log v_k^c \end{aligned}$$

We now write optimization with respect to  $v_i^a$ . We now take derivative with respect to  $v_i^a$  and set it equal to 0.

$$\begin{aligned} 0 &= -\tau^a + \sum_{jk} (z_{ijk}) \frac{1}{v_i^a} \quad (53) \\ \tau v_i &= \sum_{jk} z_{ijk} \\ v_i^a &= \frac{1}{\tau^a} \sum_{jk} z_{ijk} \end{aligned}$$

Observe that  $v_i^a \propto \sum_{jk} z_{ijk}$ . Since it is the case that  $1^t v^a = 1$ . Then the following is true.

$$\tau^a = \sum_{ijk} z_{ijk} \quad (54)$$

The optimizing updates for  $z_{ijk}$  set  $z_{ijk}$  proportional to  $\theta_{ijk} v_i^a v_j^b v_k^c$  based on the standard application of the tightest bound for Jensen's inequality. Updates for  $v_j^b$  and  $v_k^c$  follow the same form as  $v_i^a$ . Therefore the final updates are as follows.

$$\begin{aligned}
z_{ijk} &\propto \theta_{ijk} v_i^a v_j^b v_k^c \\
v_i^a &\propto \sum_{jk} z_{ijk} \\
v_j^b &\propto \sum_{ik} z_{ijk} \\
v_k^c &\propto \sum_{ij} z_{ijk}
\end{aligned} \tag{55}$$

We repeatedly update  $z$  then  $v$  until convergence.